

Accelerating convolutions on the sphere with hybrid GPU/CPU kernel splitting

P.M. Sutter*

Department of Physics, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

UPMC Univ Paris 06, UMR7095, Institut d'Astrophysique de Paris, F-75014, Paris, France

CNRS, UMR7095, Institut d'Astrophysique de Paris, F-75014, Paris, France

Center for Cosmology and Astro-Particle Physics, Ohio State University, Columbus, OH, USA

E-mail: sutter@iap.fr

Benjamin D. Wandelt

UPMC Univ Paris 06, UMR7095, Institut d'Astrophysique de Paris, F-75014, Paris, France

CNRS, UMR7095, Institut d'Astrophysique de Paris, F-75014, Paris, France

Department of Physics, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

Department of Astronomy, University of Illinois at Urbana-Champaign, Urbana, IL 61801, USA

Franz Elsner

UPMC Univ Paris 06, UMR7095, Institut d'Astrophysique de Paris, F-75014, Paris, France

CNRS, UMR7095, Institut d'Astrophysique de Paris, F-75014, Paris, France

We present a general method for accelerating by more than an order of magnitude the convolution of pixelated function on the sphere with a radially-symmetric kernel. Our method splits the kernel into a compact real-space, and a compact spherical harmonic space component that can then be convolved in parallel using an inexpensive commodity GPU and a CPU, respectively. We provide models for the computational cost of both real-space and Fourier space convolutions and an estimate for the approximation error. Using these models we can determine the optimum split that minimizes the wall clock time for the convolution while satisfying the desired error bounds. We apply this technique to the problem of simulating a cosmic microwave background sky map at the resolution typical of the high resolution maps of the cosmic microwave background anisotropies produced by the Planck space craft. For the main Planck CMB science channels we achieve a speedup of over a factor of ten, assuming an acceptable fractional rms error of order 10^{-5} in the (power spectrum of the) output map.

Big Bang, Big Data, Big Computers

September 19-21, 2012

Laboratoire Astroparticule et Cosmologie, 10 rue A. Domon et L. Duquet, 75205 Paris 13, France

*Speaker.

1. Introduction

Cosmic microwave background (CMB) experiments, such as Planck (Planck Collaboration 2011), the Atacama Cosmology Telescope (Kosowsky 2003), the South Pole Telescope (Ruhl 2004), and CMBPol (Baumann et al. 2009) promise a great wealth of cosmological and astrophysical information (Smoot 2010). The most common operation in CMB data analysis consists of convolving a real or a synthetic map with a radial kernel. Large numbers of such smoothing or filtering operations are necessary for many critical data analysis applications, such as the simulation of CMB maps (Gorski et al. 2005), map-making from multichannel maps (Tegmark 1997; Natoli et al. 2001; Stompor et al. 2001; Patanchon et al. 2008; Sutton et al. 2010), iterative calculation of inverse covariance weighted data, e.g. in the context of optimal power spectrum estimation or Wiener filtering (Wandelt et al. 2004) wavelet analysis (Hobson et al. 1999; Martinez-Gonzalez et al. 2002; Vielva et al. 2004), point-source removal (Tegmark & de Oliveira-Costa 1998; Gonzalez-Nuevo et al. 2006), and analysis of errors. The future Euclid mission (Laureijs et al. 2011b) will resolve the sky to sub-arcsecond resolution, and one technique for identifying overdensities in such a map is via convolution with a filter.

Until recently, the near-exclusive practice in the CMB community to compute radial kernel convolutions was to use the spherical convolution theorem: transformation to spherical harmonic space, multiplying the spherical harmonic coefficients with the l -space representation of the radial kernel and back-transformation to pixel space. As a consequence of the ubiquity of radial kernel convolution for data analysis on the sphere and the ready availability of software implementing the discrete forward and backward fast Spherical Harmonic Transformation (SHTs), this has become the major application for SHTs. Interest in the actual a_{lm} is relatively rare by comparison.

Graphics Processing Units (GPUs) offer a promising solution to the computational challenges posed by radial kernel convolution to current and upcoming data sets on the sphere (Brunner et al. 2007; Barsdell et al. 2010; Fluke et al. 2011) due to their low cost and high degree of parallelism. Indeed, the recent rise of cheap GPU hardware and associated extensive programming libraries have led to their use in many applications in astrophysics, such as the analysis of the Lyman- α forest (Greig et al. 2011), dust temperature calculations (Jonsson & Primack 2010), magnetohydrodynamics (Pang et al. 2010), adaptive-mesh refinement simulations (Schive et al. 2010), analysis of data from the Murchison Widefield Array (Wayth et al. 2007), volume renderings of spectral data from the Australian Square Kilometer Array Pathfinder mission (Hassan et al. 2011), and visualizations of large-scale data sets (Szalay et al. 2008).

While GPU implementations of the SHT (Hupca et al. 2010; Szydlarski et al. 2011) have only achieved modest speed-ups, Elsner & Wandelt (2011; hereafter EW11) tackled the problem of spherical convolutions for compact radial kernels by specifically designing an algorithm adapted to benefit from high degree of parallelism and memory bandwidth for compact kernels. Compared to the serial time of a highly optimized implementation of the Fast SHT algorithm, EW11 demonstrated a speed-up of up to a factor of 60 using a commodity GPU costing \$500 with the further benefit of strongly suppressing Fourier ringing artifacts. Other approaches, such as optimizing

traditional algorithms (Muciaccia et al. 1997) and using large-scale computing resources (Gheller et al. 2007), either do not scale as efficiently or do not exploit readily available hardware.

The main limitation of the method described by EW11 is that significant speed-ups can only be achieved for relatively compact kernels. While there are still many applications for such compact kernels, truncating non-compact kernels to realize large performance improvements can lead to unreasonable artefacts in the resulting convolved maps. To take advantage of the power of GPUs with kernels of arbitrary size, we must split the given kernel between a real-space portion, which will be applied using a GPU, and an ℓ -space (i.e., Fourier) portion, which will be applied using traditional CPU methods. Each portion of the full kernel will then necessarily be truncated, resulting in a small, but predictable, error. Given an upper bound for an acceptable error, we must determine the optimal splitting between real- and ℓ -space in order to achieve maximum performance.

In this work we present `scytale`¹, a tool for splitting kernels between truncated real- and ℓ -space portions, estimating the errors due to the truncations, and discovering the optimum truncations for a given kernel. We apply this tool to determine the expected speedup when splitting a given kernel between the GPU code `ARKCoS` of EW11 and the CPU code `libpsht` of Reinecke (2011). In Section 2 we discuss our strategy for splitting kernels, estimating errors, and determining the optimum truncations. We present an analysis of the errors and our optimization results in Section 3, followed by a discussion and conclusion in Section 4.

2. Splitting formalism & Optimization Strategy

We decompose a given kernel K_ℓ into truncated ℓ -space and real-space portions, which we denote as \widehat{K}_ℓ and \widehat{K}_θ , respectively. We may then construct an approximate kernel as

$$\widetilde{K}_\ell = \widehat{K}_\ell + P_{\ell\theta} \widehat{K}_\theta, \quad (2.1)$$

where $P_{\ell\theta}$ is a Legendre transformation operator. We truncate the ℓ -space kernel to a limit ℓ_{cut} and the real-space kernel to a limit θ_{cut} . Once we have the truncated kernels, we can evaluate the resulting error by taking the fractional root mean square:

$$\sigma^2 = \alpha^2 \frac{\sigma_{\text{rms}}^2}{1/4\pi \sum (2\ell+1) K_\ell^2}, \quad (2.2)$$

where the sum run from 0 to ℓ_{max} . The constant α represents any additional errors introduced by the actual convolution, such as those caused by single-precision arithmetic and inadequate kernel interpolation, and must be empirically determined. Thus, given a particular kernel, this procedure allows us to identify values of ℓ_{cut} and θ_{cut} that satisfy a given error bound.

If a particular ℓ_{cut} and θ_{cut} satisfy an error bound, we then estimate the computational cost associated with the truncated kernels. We assume the real-space portion will be solved using `ARKCoS` on a GPU, so we denote the cost as t_{ARKCoS} . Similarly, we assume the ℓ -space kernel will be solved using the standard library `libpsht` on a CPU, and hence we will denote the cost as t_{libpsht} . The cost for applying the real-space GPU kernel is

$$t_{\text{ARKCoS}} = 0.0232s \theta_{\text{cut}} + 2.428s \quad (2.3)$$

¹We take the name from the ancient cryptographic system where only rods of a precise radius could be used to decode messages.

and the cost for the ℓ -space CPU kernel is

$$t_{\text{libpsht}} = 160\text{s} \frac{\ell_{\text{cut}}^2 \ell_{\text{max}}}{4096^3}. \quad (2.4)$$

Above, θ_{cut} is in arcminutes. To determine these scalings we used an NVIDIA GeForce GTX 480 GPU and a 2.8 GHz Intel Core2 Quad CPU. Our GPU scaling is different than the study of EW11 due to updated NVIDIA drivers. Note that the CPU timing assumes the use of only a single core. We assume throughout a data set with HEALPix (Gorski et al. 2005) resolution $n_{\text{side}} = 2048$ and $\ell_{\text{max}} = 4096$, consistent with Planck observations (Mennella et al. 2011). Furthermore, we assume a power spectrum derived from WMAP 7-year results (Komatsu et al. 2011).

We assume that the GPU and CPU portions can be solved in parallel, and hence our goal for a given kernel is to find the pair $(\ell_{\text{cut}}, \theta_{\text{cut}})$ that satisfies the error bound and at which $t_{\text{ARKCoS}} = t_{\text{libpsht}}$, minimizing the overall cost.

3. Results

We study radially-symmetric kernels of the type

$$K_{\ell} = \sqrt{C_{\ell}} B_{\ell}, \quad (3.1)$$

where C_{ℓ} is the expected power in the given ℓ -space bin and B_{ℓ} is the Legendre transform of a beam. We assume an identical band limit of ℓ_{max} for both the input power spectrum and the kernel. These particular kernels have a wide variety of applications. We assume a Gaussian beam with a given FWHM. For this analysis, we will also assume $C_{\ell}^{\text{input}} \sim 1$ (that is, the case of simulating CMB maps with uncorrelated noise).

We begin with an analysis of splitting a single kernel. We show in Figure 1 an example kernel produced with a 7 arcmin FWHM beam. We truncate the kernel and the input power spectrum at $\ell_{\text{max}} = 4096$. This narrow beam produces wide support to significantly high ℓ : only past $\ell \approx 2000$ does the kernel drop below 1% of $\sqrt{C_{\ell}}$.

After splitting, the ℓ -space kernel faithfully reproduces the low- ℓ portion of the full kernel while the real-space kernel matches the high- ℓ regime. In order to fit the behavior of the full kernel at high ℓ , the real-space kernel produces large oscillations at low ℓ , which are compensated by percent-level adjustments in the ℓ -space kernel. Summed together, these kernels reproduce the full input kernel, except at the very highest ℓ where the low magnitudes make a full fit difficult.

Even though our computational approach damps oscillations in ℓ -space (where the fits to the full input kernel take place) we see rapid oscillations in the actual kernel that ARKCoS uses in its real-space approach. We must accurately interpolate this kernel, especially at small angles, in the convolution algorithm in order to both recover the high- ℓ behavior and correctly calculate the systematic offsets present in the low- ℓ portion of the approximate kernel. To do this, we employ a simple bias where we place half the available interpolation nodes within the first 1/16th of the available support; in this case, within 7.5 arcmin. We found this bias to be a good compromise between the need to carefully interpolate the innermost portions of the kernel and the need to maintain a sufficient number of interpolation points throughout the rest of the kernel.




Figure 1: Example kernel for a beam with 7 arcmin FWHM.

The approximate kernel faithfully represents the full input kernel below the truncation threshold of the ℓ -space kernel at $\ell = 1500$, which we see in Figure 2. In this figure we show the relative error, defined as

$$\sigma_\ell = \log_{10} \left| 1 - \frac{\tilde{K}_\ell}{K_\ell} \right|. \quad (3.2)$$

In this figure we see three distinct regimes. The first, from $\ell = 0$ -1500 where the ℓ -space kernel dominates, has essentially zero error. From $\ell = 1500$ to roughly 3000, we maintain a relative error of roughly 10^{-5} . In this region the real-space kernel is best able to reproduce the full input kernel. Finally, at the highest ℓ , the real-space kernel has difficulty following the input kernel and the errors begin to exponentially diverge, reaching 100% relative error at $\ell_{\max} = 4096$. However, the beam strongly suppresses the kernel here and the high-magnitude low- ℓ portion dominates our error estimate. Therefore we can ultimately satisfy a given overall error bound.

To evaluate the actual performance of each kernel, we applied them to a uniform-noise input map and extracted the spectra. We compare these spectra in Figure 3. We show the power spectrum after convolving with the full ℓ -space kernel K_ℓ , the truncated ℓ -space kernel \tilde{K}_ℓ , and the truncated real-space kernel \hat{K}_θ . We also show the power spectrum of the summed map. We see that we are able to recover the desired power spectrum using the truncated kernels, except at the highest ℓ range, where interpolation errors and the limitations of single-precision arithmetic in the GPU introduce deviations.




Figure 2: Estimated relative error of the example approximate kernel \tilde{K}_ℓ to the full kernel K_ℓ . Shown is the relative error as a function of ℓ . For this example, the ℓ -space kernel is truncated to $\ell_{\text{cut}} = 1500$ and the real-space kernel to $\theta_{\text{cut}} = 240$ arcmin.

Figure 4 shows the relative error between the power spectrum obtained by summing the maps produced by the truncated kernels and spectrum obtained by using the full ℓ -space kernel. We see similar structure to the estimated relative error, but in this case the errors are not negligible below $\ell_{\text{cut}} = 1500$. Here, the difficulty of adding the small component due to the real-space kernel to the ℓ -space kernel is apparent. After $\ell = 1500$ we see small oscillations around the full power spectrum followed by the expected exponential rise in the relative error. Altogether, we found the total error to be a factor of five higher than estimated due to these numerical effects. Thus we set the constant α in Eq.(2.2) to five.

In Figure 5 we show the map after convolving with the full ℓ -space kernel. We also show the difference between this map and sum of the maps produced by convolution with the truncated ℓ -space and real-space kernels. We maintain small errors throughout the entire map, with the largest errors at the smallest scales, as expected. In Figure 6 we show a 5-degree patch of the same maps. We see that the ℓ -space kernel reproduces the full map to percent-level accuracy. However, the real-space kernel is necessary to correctly construct the small-scale power and reduce the error to acceptable limits.

We compare our estimated RMS error to the actual map and power spectra errors in Figure 7 for a selection of ℓ_{cut} values with a fixed $\theta_{\text{cut}} = 240$ arcmin and the same 7 arcmin beam that we




Figure 3: Derived power spectra after convolving a uniform-noise map with various kernels. The kernels used are: the full ℓ -space kernel K_ℓ (pink), the truncated ℓ -space kernel \hat{K}_ℓ (red), and the truncated real-space kernel \hat{K}_θ (green). The blue line shows the power spectrum of the map created by summing the individual maps of the two truncated kernels. For this example, the ℓ -space kernel is truncated to $\ell_{\text{cut}} = 1500$ and the real-space kernel to $\theta_{\text{cut}} = 240$ arcmin.

have thus far used. For this plot, we have set the empirically-determined constant α to five. With this chosen constant, our error estimate matches the actual error in the power spectra until an ℓ_{cut} of 2500. At higher ℓ_{cut} values, we overestimate the spectrum errors, but since this lies below our chosen error bound of 10^{-5} (see below) we choose to maintain this value of α . The maps tend to produce higher errors, but since our quantity of interest is the derived power spectrum, we choose to match those errors.

With all this in place we now turn to our scanning strategy and results of our optimization study. We examine beams with 5-10 arcmin FWHM, which are most relevant to the Planck mission (Mennella et al. 2011). Table 1 shows the optimum $(\ell_{\text{cut}}, \theta_{\text{cut}})$ pairs for five of the ten beam sizes studied, assuming a maximum error bound of 10^{-5} . Below 6 arcmin we could not find suitable truncations that still maintained our desired error bound. We see that all truncations are essentially identical, indicating that the ability to split these kernels is binary: either no optimum truncations can be found, but if optimum truncations can be found they will be very aggressive. For these beam sizes, the optimum ℓ_{cut} values that satisfy the error bounds are significantly below ℓ_{max} , which promise significant enhancements in performance.

We show in Figure 8 the speedup versus beam FWHM for these beam sizes and our error




Figure 4: Actual relative error of the approximate kernel \tilde{K}_ℓ to the full kernel K_ℓ after convolution. Shown is the relative error as a function of ℓ . For this example, the ℓ -space kernel is truncated to $\ell_{\text{cut}} = 1500$ and the real-space kernel to $\theta_{\text{cut}} = 240$ arcmin.

Table 1: Optimum ℓ_{cut} and θ_{cut} pairs for each beam FWHM studied, assuming an error bound of 10^{-5} .

Beam FWHM (arcmin)	ℓ_{cut}	θ_{cut} (arcmin)
7	1158	390
8	1070	390
9	1055	360
10	979	360
11	1014	330
12	960	330
13	940	330
14	929	300
15	961	270




Figure 5: (a) Map after convolving a uniform-noise input map with the full ℓ -space kernel K_ℓ . (b) The difference between the map in panel (a) and the map constructed by summing the convolution outputs of the truncated ℓ -space kernel \hat{K}_ℓ and the truncated real-space kernel \hat{K}_θ . For this example, the ℓ -space kernel is truncated to $\ell_{\text{cut}} = 1500$ and the real-space kernel to $\theta_{\text{cut}} = 240$ arcmin.




Figure 6: (a) Five-degree patch of the map convolved with the full unsplit kernel. (b) Difference between the map in panel (a) and the map produced by convolving with the truncated ℓ -space kernel \hat{K}_ℓ . (c) Map created by convolving with the truncated real-space kernel \hat{K}_θ . For this example, the ℓ -space kernel is truncated to $\ell_{\text{cut}} = 1500$ and the real-space kernel to $\theta_{\text{cut}} = 240$ arcmin.




Figure 7: Estimated RMS error computed by *scytale* with $\alpha = 5$ (red stars) versus actual RMS error in the maps produced by convolution with a uniform-noise map (green circles) and the RMS error in the power spectra derived from those maps (blue triangles). The lines do not represent data but are shown as visual aids.




Figure 8: Speedup versus beam FWHM assuming an overall error bound of 10^{-5} .

bound of 10^{-5} . We define the scaling as the time to solution with our split approach relative to the cost of applying the entire kernel (i.e., up to ℓ_{\max}) on the CPU with `libpsht`. Below 7 arcmin, we find no optimum truncations and hence do not show them. We see significant performance gains above 7 arcmin, with the speedups plateauing in the range 12-15. This speedup implies a reduction in the computational time from 160 seconds to approximately 12 seconds for a single convolution operation. Since all the truncations are essentially the same above 7 arcmin, we find nearly identical speedups regardless of the beam size.

4. Conclusions

We have introduced and discussed a method for splitting radially-symmetric kernels into truncated real- and Fourier-space components and estimating the errors associated with such splitting. We have validated our error estimation by performing convolutions with the truncated kernels and computing the actual resulting error. We have found that for Planck-sized data sets, a large range of kernels can be split into significantly truncated portions while still maintaining an acceptable ($\sim 10^{-5}$) error bound, leading to significant speedups.

Our analysis was focused on an ideal case; i.e., situations where there is no noise and where the input power spectrum remains flat. This is the worst-case scenario. In the case where noise

dominates the high- ℓ regime we found speedups of order ~ 20 , since we could relax the criterion of strictly matching the structure of the full kernel in this region.

Our approach is currently limited to $\ell \sim 4000$ due to the finite amount of fixed memory available on single current-generation GPUs. An all-sky convolution up to $\ell = 8000$ or 16000 would require splitting the problem across multiple GPUs, as discussed below. However, current experiments that probe this regime, such as ACT (Kosowsky 2003) and SPT (Ruhl 2004), only map on the order of hundreds of square degrees. By re-orienting their survey maps onto the polar cap, we can keep the number HEALPix rings small and exploit our algorithm with currently-available GPUs.

The ARKCoS code also has a CPU-based implementation, allowing our approach to work on homogeneous architectures. While the speedups in the CPU-only case are not as significant, we can still take advantage of the parallelism offered by the compact real-space kernels. In this scenario, the truncated ℓ -space kernel can be convolved using traditional parallel spherical harmonic transform operations on a few cores (such as `ccSHT`²), where the parallel scalability is strongest, while the truncated real-space kernel can be convolved using many cores in parallel in the manner described above.

Kernel splitting enables the efficient allocation of resources for tackling large data sets; in our case, by applying real-space kernels with a GPU and ℓ -space kernels with a CPU. We have applied this kernel splitting scheme to an optimization study to find the realistic speedups associated with splitting a kernel between a compact portion to be solved on a GPU and the remainder on a CPU. Applying this to kernels and data sets appropriate for the Planck mission, we find that this splitting technique can lead to over a factor of ten speedup compared to traditional fully CPU-based approaches. This significantly improves the feasibility of many necessary and important data analysis operations, such as wavelet analysis, point source removal, and map making.

Acknowledgments

The authors acknowledge support from NSF Grants AST-0908902, AST-0908693 ARRA, and AST-0708849. This material also is based upon work supported in part by NSF Grant AST-1066293 and the hospitality of the Aspen Center for Physics.

References

- Barsdell, B. R., Barnes, D. G., & Fluke, C. J. 2010, in *Astronomical Society of the Pacific Conference Series*, Vol. 434, *Astronomical Data Analysis Software and Systems XIX*, ed. Y. Mizumoto, K.-I. Morita, & M. Ohishi, 209
- Baumann, D. et al. 2009, *Probing Inflation with CMB Polarization*, Vol. 1141 (AIP), 10
- Brunner, R. J., Kindratenko, V. V., & Myers, A. D. 2007, e-print arXiv: 0711.3414
- Efstathiou, G. 2007, *Mon. Not. Roy. Ast. Soc.*, 380, 1621

²<http://crd-legacy.lbl.gov/cmc/ccSHTlib/doc/index.html>

- Elsner, F. & Wandelt, B. D. 2011, *Astron. and Astrophys.*, 532, A35
- Fluke, C. J., Barnes, D. G., Barsdell, B. R., & Hassan, A. H. 2011, *Publications of the Astronomical Society of Australia*, 28, 15
- Gheller, C., Leach, S., Taffoni, G., Vuerli, C., & Pasian, F. 2007, *Astronomical Data Analysis Software and Systems XVI ASP Conference Series*, 376
- Gonzalez-Nuevo, J., Argüeso, F., López-Caniego, M., Toffolatti, L., Sanz, J. L., Vielva, P., & Herranz, D. 2006, *Mon. Not. Roy. Ast. Soc.*, 369, 1603
- Gorski, K. M., Hivon, E., Banday, A. J., Wandelt, B. D., Hansen, F. K., Reinecke, M., & Bartelmann, M. 2005, *Astrophys. J.*, 622, 759
- Greig, B., Bolton, J. S., & Wyithe, J. S. B. 2011, *Mon. Not. Roy. Ast. Soc.*, 1539
- Hassan, A. H., Fluke, C. J., & Barnes, D. G. 2011, in *Astronomical Society of the Pacific Conference Series*, Vol. 442, *Astronomical Data Analysis Software and Systems XX*, ed. I. N. Evans, A. Accomazzi, D. J. Mink, & A. H. Rots, 207
- Hobson, M. P., Jones, A. W., & Lasenby, A. N. 1999, *Mon. Not. Roy. Ast. Soc.*, 309, 125
- Hupca, I. O., Falcou, Â., Grigori, Â., & Stompor, Â. 2010, eprint arXiv:1010.1260
- Jonsson, P. & Primack, J. R. 2010, *New. Astron.*, 15, 509
- Komatsu, E. et al. 2011, *Astrophys. J.*, 192, 18
- Kosowsky, A. 2003, *New Astronomy Reviews*, 47, 939
- Laureijs, R., Amiaux, J., Arduini, S., Auguères, J. ., Brinchmann, J., Cole, R., Cropper, M., Dabin, C., Duvet, L., Ealet, A., & et al. 2011a, *ArXiv e-prints*
- Laureijs, R. et al. 2011b, *Euclid Definition Study Report*, arXiv: 1110.3193
- Martinez-Gonzalez, E., Gallegos, J. E., Argueso, F., Cayon, L., & Sanz, J. L. 2002, *Mon. Not. Roy. Ast. Soc.*, 336, 22
- Mennella, A. et al. 2011, eprint arXiv:1101.2038
- Miller, G. & Hoffman, C. 1984, *SIGGRAPH 84 Advanced Computer Graphics Animation seminar notes*
- Muciaccia, P. F., Natoli, P., & Vittorio, N. 1997, *Astrophys. J.*, 488, L63
- Natoli, P., de Gasperis, G., Gheller, C., & Vittorio, N. 2001, *Astronomy and Astrophysics*, 372, 346
- Pang, B., Pen, U.-l., & Perrone, M. 2010, e-print arXiv:1004.1680
- Patanchon, G. et al. 2008, *Astrophys. J.*, 681, 708

- Planck Collaboration. 2011, eprint arXiv:1101.2022
- Planck HFI Core Team. 2011, eprint arXiv:1101.2048
- Reinecke, M. 2011, *Astron. and Astrophys.*, 526, A108
- Ruhl, J. 2004, *The South Pole Telescope*, Vol. 5498 (SPIE), 11
- Schive, H.-Y., Tsai, Y.-C., & Chiueh, T. 2010, *Astrophys. J.*, 186, 457
- Smoot, G. F. 2010, in *American Institute of Physics Conference Series*, Vol. 1241, American Institute of Physics Conference Series, ed. J.-M. Alimi & A. Fuözfa, 3
- Stompor, R. et al. 2001, *Phys. Rev. D*, 65
- Sutton, D. et al. 2010, *Mon. Not. Roy. Ast. Soc.*, 407, 1387
- Szalay, T., Springel, V., & Lemson, G. 2008, e-print arXiv:0811.2055
- Szydlarski, M., Esterie, P., Falcou, J., Grigori, L., & Stompor, R. 2011
- Tegmark, M. 1997, *Astrophys. J.*, 480, L87
- Tegmark, M. & de Oliveira-Costa, A. 1998, *Astrophys. J.*, 500, L83
- Vanicek, P., Janak, J., & Featherstone, W. 2003, *Studia Geophysica et Geodaetica*, 47, 455
- Vielva, P., Martínez-González, E., Barreiro, R. B., Sanz, J. L., & Cayon, L. 2004, *Astrophys. J.*, 609, 22
- Wandelt, B., Larson, D., & Lakshminarayanan, A. 2004, *Physical Review D*, 70, 12
- Wayth, R., Dale, K., Greenhill, L. J., Mitchell, D. A., Ord, S., & Pfister, H. 2007, in *Bulletin of the American Astronomical Society*, Vol. 38, American Astronomical Society Meeting Abstracts, 744
- Yeo, B., Yu, P., Grant, P., Fischl, B., & Golland, P. 2008, *Proceedings of the International Conference on Medical Image Computing and Computer Assisted Intervention*, 5241, 468